

LDRD Final Report – New Directions for Algebraic Multigrid: Solutions for Large Scale Multiphysics Problems

V. E. Henson

February 6, 2003

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

This report has been reproduced directly from the best available copy.

Available electronically at <http://www.doc.gov/bridge>

Available for a processing fee to U.S. Department of Energy
And its contractors in paper from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

Available for the sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/ordering.htm>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

Final Report
LDRD Project
Exploratory Research in the Disciplines

**New Directions for Algebraic Multigrid:
Solutions for Large Scale Multiphysics Problems**

Tracking Code: 00-ERD-18

Three Year Budget: \$1,023,000

Van Emden Henson, *PI*
Center for Applied Scientific Computing
Computation Directorate

Principal LLNL Collaborators
Panayot S. Vassilevski
Ulrike Meier Yang
Center for Applied Scientific Computing

Subcontractors and External Collaborators
Marian Brezina, University of Colorado
Timothy Chartier, University of Washington
Todd Coffey, North Carolina State University
Craig Douglas, University of Kentucky
Kyle Gallivan, Florida State University
Johannes Kraus, Universität Leoben, Austria
Thomas A. Manteuffel, University of Colorado
Steven F. McCormick, University of Colorado
John Ruge, University of Colorado

ABSTRACT

The purpose of this research project was to investigate, design, and implement new algebraic multigrid (AMG) algorithms to enable the effective use of AMG in large-scale multiphysics simulation codes. These problems are extremely large; storage requirements and excessive run-time make direct solvers infeasible. The problems are highly ill-conditioned, so that existing iterative solvers either fail or converge very slowly. While existing AMG algorithms have been shown to be robust and stable for a large class of problems, there are certain problems of great interest to the Laboratory for which no effective algorithm existed prior to this research.

The research was primarily carried out in three main topic areas 1) the discovery and development of a family of methods based on new interpolation routines, called *element-based algebraic multigrid*; 2) innovative methods for devising coarse grids, including *element agglomeration* parallelization of advanced graph-theoretic methods; and 3) research into the application of novel smoothing methods. In the first area, the research resulted in the discovery and development of a family of methods based on a new heuristic describing an essential characteristic of AMG. These methods, AMGe, element-free and multilevel element-free AMGe, and spectral AMGe, all are based on discovering and characterizing the essential nature of smooth error components through the use of local, or neighborhood, matrices. The second area of research led to improved efficiency and effectiveness in the coarsening methods for *BoomerAMG*, the main AMG code developed by CASC, as well as to agglomeration-based methods for coarsening AMGe problems. The last area was essential to making the methods converge optimally, or, for some problems, converge at all.

The result of this research is a body of ideas, methods, algorithms, and codes that are steadily increasing both the effectiveness of CASC's solvers on multiphysics simulation problems and the range of problems to which these solvers may be applied. This research has had a demonstrable effect on the performance of at least two of the Laboratory's large-scale simulation codes. The wealth of ideas generated in this research will provide material for important further research for some time to come, research whose continuation has already found funding support from other venues.

BACKGROUND

Algebraic multigrid (AMG) has become one of the central areas of research in CASC. Largely, this is because of our customers' increasing reliance on large, unstructured grids for their multiphysics simulation packages. The sizes of many problems of interest to the Laboratory dictate that parallel methods of solution be employed, using hundreds or thousands of processors. Such massively parallel efforts can only succeed if the algorithms they rely upon are *scalable*, meaning the run-time remains constant as both the number of processors and the problem size grow in tandem.

It is well-known that for structured-grid problems multigrid methods are among the most efficient solvers available. More importantly, multigrid algorithms for many such problems are provably scalable, a theoretical fact that has been shown to be attainable in practice. Indeed, CASC has several structured-grid multigrid variants that have proven particularly useful in large-scale simulations for precisely this reason [1]. However, these methods are not applicable to unstructured-grid problems, and for those, CASC has placed much emphasis on the development of AMG algorithms.

AMG was first introduced in the early 1980's [2], and immediately attracted substantial interest [17] [18] [19]. Research continued sporadically through the late 1980's and early 1990's.

Later, there was a resurgence of interest in both AMG and related algebraic-type multilevel methods [23] [24]. This resurgence in AMG research is due to the need to solve increasingly larger systems, with hundreds of millions or billions of unknowns, on unstructured grids.

At the time this project began, CASC had developed a parallel AMG code [11] [12], and was in the process of integrating it into the *hypre* library and thence into several of the applications codes of our customers. Along the way, we had shown that AMG is both robust and algorithmically scalable for a large class of problems [6].

However, we had also discovered that there are certain classes of problems for which AMG, as currently known, is insufficient [8]. In fact, for some of these problems no scalable, efficient solver was known. For example, although AMG will work well for general elasticity problems posed on simple Cartesian meshes, such problems posed using finite elements with very high aspect ratios cause AMG (and all other sparse linear solvers) to converge very slowly. Similarly, extreme thin-body elasticity problems also cause severe difficulties. Problems involving contact surfaces and constraints are also problematic, with no efficient solvers known. All of these features, however, are of particular importance in the various simulation packages under development around the Laboratory and the need for an efficient solver for these problems is critical.

The purpose of this LDRD project, then, was to discover new AMG algorithms to address the difficulties encountered in the large-scale multiphysics simulations, to develop these algorithms into codes that served as enabling technologies, allowing the application of the new algorithms to real world problems, and, finally, to assist and facilitate CASC scientists in the integration of these new methods into Laboratory simulation codes.

TECHNICAL APPROACH

We proposed solving these difficulties by researching and developing new types of AMG. Our research fell into three broad categories: 1) the discovery and development of a family of methods based on new interpolation routines, called *element-based algebraic multigrid*; 2) innovative methods for devising coarse grids, including *element agglomeration* parallelization of advanced graph-theoretic methods; and 3) research into the application of novel smoothing methods. By far the majority of the effort was focused in area 1). These new methods require somewhat more advance information than traditional AMG, which requires only that the matrix representing the discretized differential operator be known. Specifically, these new element-based algebraic multigrid methods, known as AMGe, element-free AMGe, multilevel element-free AMGe, and spectral AMGe (denoted ρ AMGe) require that stiffness matrices for the individual finite elements be known as well. The original ideas underlying the AMGe work in this project were discovered by CASC researchers in concert with our University collaborators [4].

Multigrid Foundation

Any multigrid algorithm works by recursive application of a two-grid method, which in turn rests on several simple facts. 1) Iterative relaxation schemes reduce the error initially, but stall. 2) When the relaxation stalls, the error in the approximation is smooth. 3) Smooth error can be represented accurately on a coarse grid, with fewer gridpoints. Moreover, because there are fewer points, then we may solve for the error more efficiently on the coarse grid. A basic multigrid algorithm uses these facts by recursively applying the two-grid scheme (h represents the fine grid, H the coarse grid):

To solve $A^h u^h = f^h$, do until convergence:

1. Smoothing: $u^h \leftarrow S(u^h, A^h, f^h)$.
2. Restrict residual: $f^H = R(f^h - A^h u^h)$.
3. Solve the coarse problem: $e^H = (A^H)^{-1} f^H$.
4. Interpolate the correction: $u^h \leftarrow u^h + P e^H$.
5. Smoothing: $u^h \leftarrow S(u^h, A^h, f^h)$.

Here S is the smoothing, P is the interpolation, R is the restriction, and A^H is the coarse-grid operator. Applying recursion at step 3 produces the V -cycle, the fundamental multigrid algorithm. AMG uses precisely the same steps as just listed, and differs from conventional multigrid in that the coarse grids and associated interpolation, restriction, and coarse-grid operators are all determined automatically from algebraic considerations.

Methods Derived From Element-Based Algebraic Multigrid (AMGe).

AMGe: The first area of research was the development of AMGe, or element-based algebraic multigrid. Simply stated, AMGe is based on the discovery that by incorporating information available in the individual element stiffness matrices, it is possible to automatically determine the nature of the smooth error components. This information can then be used to create effective interpolation schemes to represent the errors, leading to effective multigrid cycling routines. Essentially, the fundamental idea underlying AMGe differs from conventional AMG in that where the conventional algorithm is guided by the heuristic that *smooth error varies slowest in the direction of dependence*, AMGe has as its guiding principle that *the interpolation operator must have the “approximation property” that it can reproduce an eigenmode of the operator matrix with error proportional to the size of the associated eigenvalue*. One reason this principle is of paramount importance in AMGe is that AMGe theory rests largely on the understanding that the eigenvectors of the operator associated with the smallest eigenvalues form a basis for the algebraically smooth error; hence it is essential that the coarse grids and interpolation operators accurately represent these important vectors. The heuristic guiding the conventional approach is grounded in purely algebraic reasoning, but is difficult to quantify accurately; the AMGe principle is derived directly from multigrid convergence theory and yields a calculable quantity, an *approximation measure* that can be used to describe the quality of the interpolation, the coarse-grid, and the accuracy of representation.

The approximation measure is given as follows [2]. Let R be an operator that “injects” a vector from the fine grid to the coarse grid and P be the interpolation operator mapping vectors from the coarse grid to the fine grid. Letting $P = QR$, we can define the *defect of interpolation* as $I - Q$. The approximation measure, M , for a given vector e , is

$$M = \frac{\langle (I - Q)e, (I - Q)e \rangle}{\langle Ae, e \rangle}.$$

Multigrid theory states that if M is bounded, over all vectors, then it can be used to determine an upper bound on the multigrid convergence factor [22]. It is easy to see that if e is an eigenvector of A associated with an extremely small eigenvalue, then the interpolation must be very accurate if M is to be bounded. Indeed, if e is in the null space of the operator, then the interpolation must be exact!

The approximation measure, however, requires global information- knowledge of the spectrum of the linear operator- which is at least as difficult to obtain as the solution to the original problem itself. An important ingredient of AMGe [4] was the development of a *local* version of the approximation measure, in terms of the spectra of the individual element stiffness matrices, and to provide a link between the global convergence theory and the local approximation measure.




Figure 1: *The elements forming the neighborhood of the node i .*

We use AMGe to create, for a given set of coarse points, an interpolation operator that is “optimal” in terms of accuracy per unit effort. We define the neighborhood of the node i , at which we wish to interpolate, as the set of points in all elements having the node i as a vertex (Figure 1). The individual stiffness matrices for all elements in the neighborhood are assembled into a *neighborhood matrix* A_i . We partition the neighborhood matrix into blocks,

$$A_i = \begin{pmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{pmatrix}$$

where A_{ff} is a block having all the interactions between the fine-grid points in the neighborhood and A_{fc} is a block containing the interactions between the fine and coarse grid points in the neighborhood. The i th row of the interpolation operator is then given by the corresponding row of

$$-(A_{ff})^{-1} A_{fc}.$$

We call this method *element interpolation*. We found that element interpolation, for certain problems (for example, 2-d thin-beam elasticity), can yield dramatic improvements in the convergence factor (the ratio of the size of the error at one step to the size of the error in the previous step), as shown in Figure 2 and Table 1.




Figure 2: The geometry of the thin-beam elasticity problem, used to demonstrate the efficacy of the AMGe element interpolation method.

Geometry		Convergence Factors	
Beam Thickness, d	Grid Spacing h	AMG Standard Interpolation	AMGe Element Interpolation
1	1/32	0.60	0.20
1/4	1/8	0.95	0.25
1/8	1/16	0.90	0.26
1/16	1/64	0.92	0.26

Table 1: A comparison of the standard interpolation and element interpolation methods on the thin-beam elasticity problem.

Element-Free AMGe: In the second year of the project we added a new area of emphasis, based on results of the first year's investigation. In the first year we determined that AMGe is indeed a more robust method than AMG, and that it may be applied to problems involving complicated grids featuring elements with bad aspect ratios and other troublesome features. Indeed, we had shown [4] that the AMGe method produces a superior interpolation operator, leading to a more robust solver. However, there remained the expense of storing and using the individual element stiffness matrices. Importantly, there are some problems (such as finite-difference codes) for which we do not have the stiffness matrices at our disposal.




Figure 3: The graph of the matrix in the vicinity of node i (red dot). The blue dots are the fine-grid neighbors, the green squares the coarse-grid neighbors, and the purple dots are the points in the exterior.

AMGe works by assembling a *neighborhood matrix*, consisting of the sum of the individual stiffness matrices for the elements having as a vertex the point at which the interpolation is desired. For the element-free method we generate the neighborhood based on the *graph of the matrix* (see Figure 3). We define three sets, the fine-grid points in the neighborhood, coarse-grid points in the neighborhood, and the *exterior*, or points that border the neighborhood. We develop the neighborhood matrix $(\hat{A}_{ff}, \hat{A}_{fc})$,

where \hat{A}_{ff} gives the fine-to-fine connections and \hat{A}_{fc} gives the fine-to-coarse connections. The neighborhood matrix is defined by the equation

$$\begin{pmatrix} \hat{A}_{ff} & \hat{A}_{fc} \end{pmatrix} = \begin{pmatrix} A_{ff} & A_{fc} & A_{fx} \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & I \\ E_{xf} & E_{xc} \end{pmatrix}$$

using E_{xf} and E_{xc} , special *harmonic extension operators* that extend vectors defined on the exterior of the neighborhood points to the fine and coarse interior points. Then, as in AMGe, we obtain the i th row of the interpolation from the i th row of $-(\hat{A}_{ff})^{-1}\hat{A}_{fc}$.

We implemented the element-free AMGe method into a serial code and tested it a number of problems [10], using several different choices for the extension operator. The results (see e.g., Table 2) indicate that on all except the most intractable problems it performs comparably to AMGe (and generally better than AMG), in terms of convergence and computational cost. Furthermore, it has the notable advantage that the storage and manipulation of the finite element stiffness matrices is not required.

Method		Elliptic	Thick Beam	Thin Beam
AMGe	iterations	23	20	19
	conv. fact.	.289	.234	.220
e-free A -extension	iterations	20	12	22
	conv. fact.	.247	.097	.280
e-free L_2 -extension	iterations	21	13	27
	conv. fact.	.231	.104	.254

Table 2. *Selected results comparing the AMGe and element-free AMGe methods. Methods are used as a preconditioner to conjugate gradient. Shown are the number of PCG iterations to achieve a specified tolerance on the residual norm and the final convergence factors.*

Multilevel Element-free AMGe: A fundamental observation to be made about AMGe is that a local calculation is performed on a small matrix (that is, an eigenvector decomposition of the neighborhood matrix), the purpose of which is to capture locally the global character of algebraic smoothness. Element-free AMGe mimics this by constructing local matrices artificially based on the graph of the matrix. Actually capturing the global character of algebraic smoothness would be preferable, as it is the global problem that is to be ultimately solved; however, capturing the global character of the smoothness entails an eigenvalue decomposition of the global operator, requiring a computational effort at least equal to that of solving the original problem, and is clearly not feasible. One branch of our research was to investigate whether the global characterization of the space of smooth vectors could be computed *locally* using a multilevel approach; that is, to compute local matrices on coarser and coarser grids and combine the spectral vectors on all levels, in much the same spirit as the multigrid algorithm itself proceeds. As a result of these investigations, the *Multilevel Element-free AMGe* method was devised. The method was implemented and tested on several problems, and yielded very promising results.

The implementation is not computationally efficient, and therefore impractical in its current form; nonetheless, it remains an active area of research.

Spectral AMGe: Every AMG algorithm depends on some concept of “smoothness.” For example, classical AMG algorithms define smooth error e as being characterized by small residual, $Ae \approx 0$. The AMGe family of methods is based on the concept that smooth error is characterized by small energy, $\langle Ae, e \rangle \approx 0$. Either of these assumptions work for many problem types; both break down in some situations. Early in our research, we concentrated on discovering

$a=1$ $b=1000$ $c=0$	$a=1$ $b=1$ $c=2$
$a=1$ $b=1$ $c=0$	$a=1000$ $b=1$ $c=2$




Figure 4: Algebraically smooth error for the problem $-(au_x)_x - (bu_y)_y - cu_{xy} = f$ where a , b , and c vary on the unit square as shown at left.

an accurate characterization for the “smooth error” that must be eliminated by coarse-grid correction. In geometric multigrid this is relatively straightforward; smoothness implies that the error is characterized by slow oscillation. In the AMG setting, however, where grids are often unstructured and/or operators are complicated, we cannot rely on a physical “smoothness.” In this case, smooth error is *defined* as the error not eliminated by relaxation, and can only be ferreted out by its algebraic properties. Indeed, for some problems, algebraic smoothness does not appear physically smooth at all (see Figure 4).

In the second year of this research we discovered a method, known as *spectral AMGe* (and denoted ρ AMGe), based on the principle that smooth error corresponds to the eigenvectors associated with the smallest eigenvalues of the operator matrix A . This fact has long been known, but characterizing the global eigenvectors of the operator matrix is a problem at least as difficult as solving the original problem $Ax = f$. The new method method, based on a “local” characterization of smoothness, uses the eigenvectors corresponding to the small eigenvalues of the neighborhood matrix (assembled from the element stiffness matrices for the agglomerated elements).

One important feature of this method is that we eliminate the requirement that the coarse-grid variables be subsets of the fine-grid variables; in this method the coarse-grid degrees of freedom

are coefficients of the eigenvectors themselves, which form the columns of the interpolation matrix. In this, spectral AMGe is more “algebraic” than the original AMGe. More importantly, this method relies on a simpler, more primitive characterization of smoothness, hence one that is less likely to break down. To date the spectral AMGe method, although expensive computationally, has proven to be one of the most robust AMG algorithms currently known.

Parallel Coarse-grid Selection and Element Agglomeration.

Element Agglomeration: One of challenges of the element-based methods (AMGe, ρ AMGe) is that even if the coarse gridpoints are known, we still must have coarse-grid “elements” and their associated stiffness matrices, in order to generate a multilevel algorithm. Hence, we need an automatic way of determining the coarse elements. We devised such a method. A form of *element agglomeration*, it creates the coarse-grid information in a two-step process. First, collecting the fine-grid elements having non-trivial intersections into agglomeration sets creates the agglomerated elements. Next, faces, edges, and vertices of the agglomerated elements are defined by algebraic abstraction of the properties of geometric elements. Specifically, to create the coarse-grid information two steps are proposed:

1. Create the set of agglomerated elements $\{E\}$ by collecting sets of fine-grid elements having non-trivial intersections into agglomeration sets.
2. Define faces, edges, and vertices of the agglomerated elements by abstracting what happens with geometric elements, as follows:
 - a. Consider all intersections $E_i \cap E_j$ for all pairs (i, j) of agglomerated elements. An intersection of this type is a *face*, F , if it is not contained in any other intersection, that is, for any pair of indices other than (i, j) .
 - b. Consider all intersections of all pairs of faces $F_i \cap F_j$. Such an intersection is called an *edge*, Γ , if it is not contained in any other intersection of faces.
 - c. Consider all intersections of all pairs of edges, $\Gamma_i \cap \Gamma_j$. Those intersections that are nonempty are called *vertices*.

Armed with the elements, faces, edges, and vertices, the local interpolation operators can be constructed by application of an energy minimization principle [14]. The global interpolation operators can then be constructed from the local operators by an assembly process not unlike standard finite-element assembly into the operator matrices. That this is both possible and effective in certain special cases (where the appropriate coarsening and agglomeration is known, such as semi-coarsening for strongly anisotropic operators) has been shown in preliminary tests by Vassilevski and Jones [14]. Abstracting the ideas to the general situation, for highly unstructured grids, or problems where the gridpoint locations are themselves unknown (such as geodetic survey problems), and then applying it in an efficient code, was among the main goals of this project

The algorithm was implemented in a sequential code and applied to general situations, e.g., for highly unstructured grids. An example of a sequence of coarse grids generated by this method is given in Figure 5. For each grid level, the individual elements are displayed as colored regions,

with different colors indicating different elements. This small problem (1600 elements on the fine grid) is illustrative; we have applied the method to much larger problems.




Figure 5. *The sequence of coarse grids automatically selected by the element agglomeration method. The finest grid (1600 elements) is at top left; the coarsest (1 element) is at bottom right. The fine-grid elements are displayed on all grids for comparative purposes.*

Parallel Coarse-grid Selection: As important as knowing *how* to agglomerate elements to produce the coarse-grid elements, however, is the problem of determining *which* points to choose as coarse-grid points and which elements to agglomerate. Although strongly related, the problem of finding a good coarse-grid can be independent of agglomeration. We needed to develop effective parallel algorithms for selecting coarse-grid points, and then forming the coarse-grid elements based on these points. For *BoomerAMG*, we had implemented an algorithm based on the parallel independent-set algorithms of Luby, Jones, and Plassman [9, 15, 16], and found that while they are efficient from a parallel processing standpoint, they tended to result in poor coarsenings. Much of our research [11, 12] focused largely on creating hybrid algorithms combining the features of the Luby-based method with quasi-parallel implementations of the standard AMG coarsening of [18]. This work resulted in considerable improvement of the implementation.

For AMGe, we showed that we can use the element information to help select the coarse-grid points. For example, the local approximation measure M can be used to determine which subset of coarse-grid points can best be used to interpolate a given fine-grid variable. We have experimented with exhaustive-search methods using this principle [5], and have developed good, robust interpolations, however, these come at an unacceptably high computational cost. While this approach has not yet proven to be practical, further research may enable us to devise a fast, effective algorithm for selecting coarse-grid points and/or elements based on the local approximation measure M .

Compatible Relaxation Coarsening: In a recent paper Brandt [3] elucidated a useful mechanism for determining the quality of a set of coarse-grid points. In essence, the idea is that relaxation is initially performed only on the F -points, holding the values at the C -points fixed, and using a homogeneous right-hand side. This relaxation method is called *compatible relaxation*. Brandt showed that those regions of the grid for which compatible relaxation converges rapidly are adequately represented on the coarse grid. Our research showed that, for element-based problems, compatible relaxation converges rapidly if and only if the AMGe measure M is small, giving us a strong relationship between the method of compatible relaxation, and its associated theory, and the theory and practice of AMGe and ρ AMGe. We have adapted the idea of compatible relaxation into the basis of a coarsening scheme, which appears as follows.

Compatible relaxation-based coarsening algorithm:

1. Initialize $U = \Omega$, the set of all points on the grid. Set $C = \emptyset$. Choose a threshold value θ .
2. While $U \neq \emptyset$, do
 - a. Perform ν sweeps of compatible relaxation
 - b. Set $U = \left\{ i : \frac{x^\nu}{x^{\nu-1}} > \theta \right\}$
 - c. Choose $S = \{\text{an independent set of points in } U\}$
 - d. Set $C = C \cup S$, $F = \Omega - C$

We implemented this algorithm in several test codes, as well as the *BoomerAMG* code, and conducted “proof of principal” tests for a limited set of problems.

Research into Smoothing Methods:

Classical AMG methods are generally based on Gauss-Seidel smoothing, in which each equation of the system is solved for the unknown associated with it, holding all other unknowns fixed at their most recent value. This method is inherently sequential, and therefore not amenable to parallelization. Another alternative is the Jacobi method, which follows the same premise except that new values for all unknowns are computed before any of these new values are applied. Jacobi, unlike Gauss-Seidel, is readily parallelizable. However, Jacobi-based methods are provably much slower than those using Gauss-Seidel; worse, our research demonstrated that as the problem size increases, the Jacobi-based method slows down, rendering it unscalable.

We expended a significant effort in the research of smoothing techniques in an effort to overcome these problems. Our standard method is a hybrid approach, that utilizes Gauss-Seidel iterations on the interior of each processor domain and a Jacobi-like parallel update on the processor boundaries. While performing better than simple Jacobi, this method is not robust, leading us to carefully examine alternatives, such as specially weighted hybrid methods and parallel multiplicative and additive Schwarz smoothers.

RESEARCH RESULTS

Over the three years of this project, we made substantial progress on the research topics described above. The nature of this project was to produce algorithms and prototypes; it was never intended that finished production codes would result. Rather, the idea was to prove up the concepts and discoveries into testable codes that could then be developed using other funding. By that measure we have been quite successful. Some of the ideas originally discovered in the course of the collaborative effort of this project have been developed, completed, and implemented in production code under the aegis of the Scalable Linear Solvers project, using ASCI, DNT, or SciDAC funding. Examples of this include the hybrid Jacobi/Gauss-Seidel and Schwarz smoothers, as well as the hybrid classical-AMG/graph-theoretic and compatible relaxation based coarsening schemes. Many of the other results were implemented in codes and tested on suites of problems, but are not yet implemented in production codes. Examples of this include the element-free and multilevel element-free AMGe methods and the spectral AMGe algorithm. These methods and codes are currently undergoing parallelization, development, and implementation into production codes, as appropriate, by the Scalable Linear Solvers team.

Still, it is important to note that some of the results of this research have had a significant and positive impact on several of the Laboratory simulation codes. Many of the improvements to the *BoomerAMG* code in CASC's *hypr* library of preconditioners are direct results of this research, for example the smoothing and coarsening methods described above. The *hypr* library has been implemented successfully in the ARIES, ALE3D, and Kull simulation codes, among others. It has been used successfully on problems involving well over a thousand processors and several millions of unknowns.

Below is a list of some of the accomplishments attained either completely or in significant part as a result of this research:

- We implemented a multilevel element-agglomeration code, in C, using a data structure for storing element stiffness matrices that maps easily to the Finite Element Interface (the FEI was developed as part of the ASCI Problem Solving Environment effort).
- We devised and implemented an algorithm that selects the coarse-grid by applying an “approximation measure with penalty” approach.
- We devised and implemented an efficient, improved parallel coarse-grid selection for *BoomerAMG*, to replace the original algorithm. This resulted in several algorithms, including both parallel implementations of the classic Ruge-Stüben algorithm as well as the so-called *Falgout* algorithm, a hybrid of the parallel Cleary-Luby-Jones-Plassman and Ruge-Stüben algorithms.
- The element-free AMGe approach was discovered. This was a classic example of serendipitous discovery; it was not one of our original goals, but arose during research discussions with our CASC and university colleagues. After its discovery, the perfection and theoretical background of this approach became an important part of the research effort of this project.
- We implemented the full multilevel element-free AMGe method in a serial code. This code is scheduled for parallelization at a later date.

- We implemented the element-agglomeration based AMGe method into a full multilevel serial code (this was an original milestone for FY2001, but was completed early, in FY2000). This code is currently being parallelized as part of the SciDAC effort.
- We integrated *BoomerAMG* into two of the Laboratory’s major parallel simulation codes (ALE3D and Kull). It remains in common use in those codes, and has been described as “the workhorse preconditioner” by a member of the Kull code group.
- We created a “systems” version of *BoomerAMG*, suitable for handling problems with coupled systems of PDEs (this capability has also been built into AMGe and element-free AMGe).
- We developed a test suite of problems for developing AMGe codes. This suite encompasses various PDE’s (particularly diffusion and elasticity problems) on various grids- both structured and unstructured.
- We determined the “best” element agglomeration strategy. A robust agglomeration algorithm was been implemented.
- We devised and implemented several improved smoothers for the *BoomerAMG* code, including hybrid Jacobi/Gauss-Seidel and overlapping-Schwarz domain-decomposition-like smoothers. These were implemented and shown to be robust for certain difficult problems.
- We discovered the ρ AMGe approach. This, like element-free AMGe, was a serendipitous discovery made in collaboration with the CASC-University of Colorado AMG research team. It became a major component of the research effort in the third year of the project.
- Leveraging the efforts of the CASC Nonlinear Solvers and Differential Equations Project (under the direction of Dr. Carol S. Woodward), we cooperatively devised and implemented the element-agglomeration based AMGe method into a nonlinear solver (for a nonlinear diffusion problem) based on the Full Approximation Scheme (FAS).

IMPACT OF THIS RESEARCH PROJECT ON THE LABORATORY.

The driving force behind this proposal was the need to develop new, efficient solvers for real-world simulations featuring the thin-body elasticity, constraint equations, slide-surfaces, high-aspect ratio elements, or highly unstructured grids that render more conventional techniques essentially ineffective. This research is at the leading edge of AMG technology, and therefore enjoys a very high profile both at the Laboratory and in the larger community of experts in linear solvers. It has immediate applicability to high-profile codes. The project has solidified the status of CASC, and therefore LLNL, as a leader in computational mathematics and computational simulations for physics.

As noted above, some of the results of our research has been implemented into the Laboratory codes ALE3D and Kull. The parallel AMG code *BoomerAMG*, incorporating numerous algorithms developed in this project, has been described as a “workhorse” preconditioner in the Kull code. We have demonstrated scalability of the Kull code, with *BoomerAMG* preconditioned

GMRES as the linear solver, on a structured-grid diffusion problem with 9.2 million unknowns on 1500 processors of the ASCI Blue machine, for similarly sized problems on the Frost machine, and for an unstructured-grid 2.3 million zone problem using 1500 processors. *BoomerAMG* is also installed in codes for gravitational potential calculations in unstructured-grid astrophysical simulations.

BoomerAMG has also been integrated into the ALE3D code. At the end of the second year of this work, we conducted tests on certain ALE3D problems that demonstrate a strong potential for success. For example, on a small spherical-shell problem (35K unknowns) on four processors, we can compare the number of iterations and time required to solve the problem using conjugate-gradient (CG), *BoomerAMG*, and preconditioned CG solvers (where we compare *BoomerAMG* and ParaSails as the preconditioner).

Solver	Iterations	Time (sec)
CG	1238	99
BoomerAMG	178	196
BoomerAMG-CG	35	44
ParaSails-CG	770	113

Table 3. Comparison of iteration counts and solution times (in seconds) of four methods on an ALE3D test matrix.

It is evident from Table 3 that *BoomerAMG*, used as a preconditioner to CG, is more effective than the other methods tested for solving this problem.

Since that time, *BoomerAMG* has been tested more fully in ALE3D. It has been successfully been used for production-type problems up to about 2 million zones in size (about six million unknowns). Problems larger than that had remained problematic through the end of the project. Quite recently, however, in the follow-on research of the SciDAC project, one of the new hybrid smoothers developed in this project has been employed. Findings remain very preliminary, but there is reason to believe that with appropriate tuning of the smoother parameters *BoomerAMG* is successful on problems significantly larger than before (we have one successful test using some 4 million zones and just over 12 million unknowns), holding out the promise that we may be able to successfully apply it to some of the truly intractable problems facing the ALE3D code group.

We also conducted tests of a more generic nature, to demonstrate the scalability and effectiveness of our research for various classes of problems. For example, we considered the family of nonsymmetric variable-convection problems $-u_{xx} - u_{yy} - u_{zz} + c(u_x + u_y + u_z) = f$ where the parameter c varies throughout the domain. There are 64K unknowns per processor on these structured-grid problems. The problem is grown by refinement in each of the three coordinate directions; hence, the number of processors grows as the cube of successive integers. In Table 4 we compare the solution times (in seconds) and number of iterations required for the *BoomerAMG*, GMRES, and BiCGSTAB methods. As is evident from the results in the table, our research has produced algorithms that are more effective, and more scalable, than other methods on these types of problems.

Processors		1	8	27	64	125	216	343	512
<i>BoomerAMG</i>	Time	8	19	23	28	34	37	46	57
	# Iter.	5	7	7	9	8	9	8	9
GMRES	Time	25	60	92	135	168	210	263	359
	# Iter.	143	273	398	534	642	793	842	1046
BiCGSTAB	Time	17	41	67	100	117	188	187	317
	# Iter.	77	144	226	323	367	493	500	644

Table 4. Comparison of *BoomerAMG*, *GMRES*, and *BiCGSTAB*, giving time to solution (in seconds) and iteration counts for each method.

TECHNICAL OUTPUT OF THE PROJECT

Publicly communicated results of our research effort include the following papers and presentations. Sixteen of the invited presentations were delivered in Europe, underscoring the fact that this project has placed CASC in the forefront of worldwide research into algebraic multilevel methods.

Papers:

- Chartier, T., Falgout, R. D., Henson, V. E., Jones, J. E., Manteuffel, T. A., McCormick, S. F., Ruge, J. W., and P. S. Vassilevski, *Spectral Agglomeration AMGe*, in preparation.
- Henson, V. E., Kraus, J. K., and Vassilevski, P. S., *Computing Interpolation Weights in AMG based on Multilevel Schur Complements*, in preparation.
- Chartier, T., Falgout, R. D., Henson, V. E., Jones, J. E., Manteuffel, T. A., McCormick, S. F., Ruge, J. W., and P. S. Vassilevski, *Spectral AMGe (ρ AMGe)*, to appear in SIAM Journal on Scientific Computing, accepted 2002.
- Henson, V. E., and U. M. Yang, *BoomerAMG: a Parallel Algebraic Multigrid Solver and Preconditioner*, Applied Numerical Mathematics, vol. 41, pp. 155-177, 2002.
- Jones, J. E., and P. S. Vassilevski, *AMGe based on element agglomeration*, SIAM Journal on Scientific Computing, vol. 23, pp. 109-133, 2001.
- Henson, V. E., and P. S. Vassilevski, *Element-free AMGe: General algorithms for computing interpolation weights*, SIAM Journal on Scientific Computing, vol. 23, pp. 629-650, 2001.
- Brezina, M., A. J. Cleary, R. D. Falgout, V. E. Henson, J. E. Jones, T. A. Manteuffel, S. F. McCormick, and J. W. Ruge, *Algebraic multigrid based on element interpolation (AMGe)*, SIAM Journal on Scientific Computing, Vol. 22, No. 5, pp. 1570-1592, 2000.
- Cleary, A. J., R. D. Falgout, V. E. Henson, J. E. Jones, T. A. Manteuffel, S. F. McCormick, G. N. Miranda, and J. W. Ruge, *Robustness and scalability of algebraic multigrid*, SIAM Journal on Scientific Computing, vol. 21, no. 5, pp 1886-1908, April, 2000.

- Brezina, M., A. J. Cleary, R. D. Falgout, V. E. Henson, J. E. Jones, T. A. Manteuffel, S. F. McCormick, and J. W. Ruge, *Algebraic multigrid based on element interpolation (AMGe)*, SIAM Journal on Scientific Computing, vol. 22, no. 5, pp. 1570-1592, 2000.

Presentations (speaker on multiauthored presentations indicated by *):

- Koley, T., Pasciak, J., and *Vassilevski, P. S., *Algebraic construction of mortar multiplier spaces with application to parallel AMGe*, Seventh European Multigrid Conference, Hohenwart Forum, October 7--10, 2002
- Vassilevski, P. S., *Spectral Agglomerate AMGe for Unstructured Finite Element Elliptic Problems*, invited plenary talk, Conference on Challenges in Scientific Computing, CISC 2002, Berlin, October 2--5, 2002.
- Vassilevski, P. S., *Algebraic Unstructured Finite Element Computations*, invited plenary talk, Numerical Methods and Applications, NMA'02, Borovets, Bulgaria, August 21, 2002
- Henson, V. E., *The State of Algebraic multigrid Research: where did we come from, where are we now, and where are we going?* invited plenary talk, Householder Symposium XV, Peebles, Scotland, June 19, 2002
- Henson, V. E., *Multigrid and Algebraic Multigrid Methods for Solving PDEs Numerically: an Overview*, invited seminar, Department of Petroleum Engineering, Stanford University, June 4, 2002
- Henson, V. E., *Algebraic Multigrid Research: why we are where we are, and where (the heck) are we going?* invited seminar talk, Pennsylvania State University, November 9, 2001
- Vassilevski, P. S., *Sparse matrix element topology with application to algebraic multigrid*, Invited presentation, PRISM 2001: Conference on Preconditioned Robust Iterative Solution Methods for Problems with Singularities, University of Nijmegen, Nijmegen, The Netherlands, May 21-13, 2001
- Henson, V. E., *Algebraic multigrid preconditioning: who, what, when, how, and why*, invited plenary presentation, 2001 International Conference On Preconditioning Techniques For Large Sparse Matrix Problems In Industrial Applications, Tahoe City, CA, April 29-May 1, 2001
- Falgout, R. D., *New Advances in Algebraic Multigrid*,
 - invited seminar, Purdue University, Indiana, April 17, 2002
 - invited seminar, University of Colorado, Boulder, Colorado, September 11, 2001
 - invited talk, Johannes Kepler University, Linz, Austria, May 25, 2001
 - invited seminar, Universität Bonn, Bonn, Germany, May 22, 2001
 - invited talk, German national Institute for Information Technology (GMD), Sankt Augustin, Germany, May 21, 2001
 - invited talk, Oberwolfach Conference, Oberwolfach, Germany, May, 2001

- *Jones, J. E., Vassilevski, P. S., and Woodward, C. S., *Computational Issues in the Application of Nonlinear Multigrid to Nonlinear Diffusion Problems*, Tenth Copper Mountain Conference on Multigrid Methods, Copper Mountain, CO, April 1-6, 2001
- Yang, U. M., *On the use of Schwarz smoothing in AMG*, Tenth Copper Mountain Conference on Multigrid Methods, Copper Mountain, CO, April 1-6, 2001.
- Vassilevski, P. S., *Spectral-element agglomerate coarsening in AMGe*, invited seminar, Texas A & M University, College Station, TX, December 6, 2000
- Vassilevski, P. S., *Spectral-element agglomerate coarsening in AMGe*, invited seminar, University of Colorado at Boulder, Boulder, CO, November 7, 2000
- Yang, U. M., *BoomerAMG: a Parallel Algebraic Multigrid Solver*, invited seminar at the GMD (German National Research Center), St. Augustin, Germany, August 28, 2000
- Henson, V. E., and U. M. *Yang, *Experiences with BoomerAMG: a Parallel Algebraic Multigrid Solver and Preconditioner for Large Linear Systems*, invited presentation, 16th IMACS World Congress on Scientific Computation, Modeling and Applied Mathematics, Lausanne, Switzerland, August 21-25, 2000
- Yang, U. M., *Parallel Algebraic Multigrid Methods*,
 - invited seminar, Central Institute of Applied Mathematics , Research Centre Juelich (FZJ), Germany, August 29, 2000
 - invited seminar, Bonn University, Bonn, Germany, August 18, 2000
- Henson, V. E., and Vassilevski, P. S., *Element-Free AMGe: General Algorithms for Computing Interpolation Weights*,
 - invited presentation, YES 2000: Yosemite Education Symposium on Advanced Multiscale and Multiresolution Methods, Yosemite National Park, CA, October 30-November 1, 2000.
 - invited seminar, University of Heidelberg, Heidelberg, Germany, July 4, 2000
 - invited seminar, Universität Bonn, Bonn, Germany, June 30, 2000
 - invited seminar, Workshop on Algebraic Multigrid, Sankt Wolfgang, Austria, June 27, 2000
- Falgout, R. D., *Parallel Algebraic Multigrid for Finite Elements (AMGe)*, invited presentation, Seminar, Pennsylvania State University, February 18, 2000
- Vassilevski, P.S., *AMGe (Algebraic multigrid) for finite element problems*,
 - invited presentation, Advanced Summer School Workshop, Pennsylvania State University, June 20, 2000
 - invited presentation, Colloquium talk, University of California at San Diego, February 17, 2000
 - invited presentation, Numerical Analysis Seminar, Texas A&M University, College Station, TX, January 12, 2000

THE PRINCIPAL INVESTIGATOR AND THE RESEARCH TEAM

Principal Investigator. **Dr. Van Emden Henson** is a computational Mathematician at CASC and the Group Leader for the Scalable Algorithms Group. He is recognized as an expert in multigrid methods, and has concentrated on AMG in recent years. He is one of the original developers of *BoomerAMG*. Dr. Henson's research interests include multigrid methods, image reconstruction from projections, and discrete Fourier transforms. He is a coauthor of two books, *The DFT: an owner's manual for the discrete Fourier transform*, and *A Multigrid Tutorial, 2nd Edition*. Prior to joining CASC Dr. Henson spent seven years as an Assistant Professor of Mathematics at the Naval Postgraduate School in Monterey, California. He earned his PhD in Applied Mathematics from the University of Colorado at Denver in 1990, following a decade as an Exploration Geophysicist with Cities Service Oil & Gas Corporation and Occidental Petroleum.

Co-investigators. **Dr. Panayot Vassilevski**, a computational mathematician, joined CASC in 1998. He earned his PhD in Mathematics from University of Sofia, Bulgaria, in 1984. He is an expert on numerical PDEs, both discretization and solution methods, particularly multilevel iterative methods, hierarchical basis and multigrid, and domain decomposition methods. He also has expertise in more traditional preconditioning methods such as incomplete (block) factorization methods. Prior to joining CASC/LLNL he was an Associate Professor at the Bulgarian Academy of Sciences, and has been a visiting Associate Professor at several Universities: UCLA (1991-1993), Texas A & M (Spring 1996), Bowling Green State University (Spring 1997), and the University of California at San Diego (Winter 1998). He is a member of the editorial board of the *SIAM Journal of Numerical Analysis* and Associate Editor of the *Journal of Numerical Linear Algebra with Applications*. **Dr. Ulrike Meier Yang** is a computer scientist who joined CASC in 1998. She earned her PhD in Computer Science from the University of Illinois at Urbana-Champaign in 1995. Dr. Yang is an expert on supercomputing, parallel computing, Krylov methods, and preconditioned iterative solvers. Prior to joining CASC she worked as a Senior Computer Scientist at the Center for Supercomputing Research and Development at the University of Illinois at Urbana-Champaign and also as a Computer Scientist at KFAJuelich GmbH, in Germany.

In-House Collaborations. CASC has the nation's largest group of multigrid researchers. Several served in a collaborative role in this project. In particular, we relied on advice, ideas, and assistance from **Dr. Jim E. Jones** on AMGe and element-agglomeration topics, **Dr. Robert D. Falgout** on parallel multigrid and parallel AMG techniques, **Dr. Charles Tong** on smoothed aggregation algebraic multigrid and Krylov methods, and **Dr. Andrew Cleary** on parallel coarsening methods. Drs. Tong, **Dr. Edmond Chow** and **Dr. David Stevens** were instrumental as liaisons to the ALE3D code group, while **Dr. Michael Lambert** performed a similar function for the Kull code group.

External Collaborations. We maintained several strategic, long-term alliances with colleagues in academia, a number of whom contributed to this research effort. Most important among these is our alliance with researchers from the University of Colorado, including **Dr. Thomas A. Manteuffel**, **Dr. Stephen F. McCormick**, **Dr. Marian Brezina**, and **Dr. Timothy Chartier** (now at the University of Washington), as well as **Dr. John W. Ruge** of Front Range Scientific Computations, Inc. All of these scientists are experts in multigrid and AMG; indeed, Drs.

McCormick and Ruge are among the founders of that field. In addition, we established collaborations with **Dr. Craig C. Douglas**, of the University of Kentucky, and **Dr. Kyle Gallivan**, of Florida State University. Dr. Douglas is a well-known expert in multigrid methods and in advanced scientific computing, and is emphasizing cache-based algorithms. Dr. Gallivan is an expert on performance analysis and enhancement of high-performance computers.

Subcontracts. Some portions of this research were performed by external contractors. Specifically, subcontracts funded by this research project were awarded to Dr. John Ruge of Front Range Scientific Computations, Inc., and Dr. Marian Brezina of the University of Colorado, who together examined some details regarding the implementation and effectiveness of AMGe, Dr. Timothy Chartier of the University of Washington, who implemented and extended our research into ρ AMGe, and Dr. Todd Coffey of North Carolina State University, who examined the use of AMGe and element agglomeration for nonlinear problems.

REFERENCES

- [1] Baldwin, C., P. N. Brown, R. D. Falgout, J. Jones, and F. Graziani, *Iterative linear solvers in a 2d radiation-hydrodynamics code: methods and performance*, Journal of Computational Physics, 154, 1999, pp. 1-40.
- [2] Brandt, A., *Algebraic multigrid theory: The symmetric case*, Appl. Math. Comput., 19, 1986, pp. 23-56.
- [3] Brandt, A., *General Highly Accurate Coarsenings*, Electronic Transactions in Numerical Analysis, vol. 10, pp. 1-20, 2000.
- [4] Brezina, M., A. J. Cleary, R. D. Falgout, V. E. Henson, J. E. Jones, T. A. Manteuffel, S. F. McCormick, and J. W. Ruge, *Algebraic multigrid based on element interpolation (AMGe)*, SIAM Journal on Scientific Computing, Vol. 22, No. 5, pp. 1570-1592, 2000.
- [5] Chartier, T., *Algebraic multigrid based on element interpolation (AMGe): Coarsening and the AMGe measure*. Presented at the 10th Copper Mountain Conference on Multigrid Methods, April 1999.
- [6] Cleary, A. J., R. D. Falgout, V. E. Henson, J. E. Jones, T. A. Manteuffel, S. F. McCormick, G. N. Miranda, and J. W. Ruge, *Robustness and scalability of algebraic multigrid*, SIAM Journal on Scientific Computing, vol. 21, no. 5, pp. 1886-1908, April 2000.
- [7] Douglas, C. C., J. Hu, M. Kowarschik, U. R de, and C. Weiss, *How much can be gained by cache-aware programming of multigrid methods*, presented at the 10th Copper Mountain Conference on Multigrid Methods, April 1999.
- [8] Falgout, R. D., V. E. Henson, J. E. Jones, and U. M. Yang, *Some recent advances in algebraic multigrid*, presented at the 10th GAMM Workshop on Multigrid Methods, University of Bonn, Germany, October 1998.
- [9] Gjertsen, R. K., Jr., M. T. Jones, and P. E. Plassman, *Parallel heuristics for improved, balanced graph colorings*, Journal of Parallel and Distributed Computing, 37, 1996, pp. 171-186.
- [10] Henson, V. E., and P. S. Vassilevski, *Element-free AMGe: General algorithms for computing interpolation weights*, SIAM Journal on Scientific Computing, vol. 23, pp. 629-650, 2001.

- [11] Henson, V. E., and U. M. Yang, *Experiences with BoomerAMG: A Parallel Algebraic Multigrid Solver*, Applied Numerical Mathematics, vol. 41, pp. 155-177, 2002.
- [12] Henson, V. E., and U. M. Yang, *BoomerAMG: A parallel implementation of algebraic multigrid*, presented at the 9th SIAM Conference on Parallel Processing, March 1999.
- [13] Hu, J., *Cache-based multigrid on unstructured grids*, presented at the 10th Copper Mountain Conference on Multigrid Methods, April 1999.
- [14] Jones, J. E., and P. S. Vassilevski, *AMGe based on element agglomeration*, SIAM Journal on Scientific Computing, vol. 23, pp. 109-133, 2001.
- [15] Jones, M. T., and P. E. Plassman, *A parallel graph coloring heuristic*, SIAM Journal on Scientific Computing, 14, 1993, pp. 654-669.
- [16] Luby, M., *A simple parallel algorithm for the maximal independent set problem*, SIAM Journal on Computing, 15, 1986, pp. 1036-1053.
- [17] Ruge, J. W., and K. Stüben, *Efficient solution of finite difference and finite element equations by algebraic multigrid (AMG)*, in Multigrid Methods for Integral and Differential Equations, D. J. Paddon and H. Holstein, eds., The Institute of Mathematics and its Applications Conference Series, Clarendon Press, Oxford, 1985, pp. 169-212.
- [18] ———, *Algebraic multigrid (AMG)*, in Multigrid Methods, S. F. McCormick, ed., vol. 3 of Frontiers in Applied Mathematics, SIAM, Philadelphia, PA, 1987, pp. 73-130.
- [19] Stüben, K., *Algebraic multigrid (AMG): experiences and comparisons*, Appl. Math. Comput., 13, 1983, pp. 419-452.
- [20] Stüben, K., *Algebraic Multigrid (AMG): An Introduction with Applications*, GMD Technical Report 53, GMD-Forschungszentrum Informationstechnik GmbH, Sankt Augustin, Germany, 1999.
- [21] Vaněk, P., M. Brezina, and R. Tezaur, *Two-grid method for linear elasticity on unstructured meshes*, SIAM J. Sci. Comput., 1998.
- [22] ———, *Algebraic multigrid based on smoothed aggregation for second and fourth order problems*, Computing, 56 (1996), pp. 179-196.
- [23] Vaněk, P., J. Mandel, and M. Brezina, *Algebraic multigrid on unstructured meshes*, UCD/CCM Report 34, Center for Computational Mathematics, University of Colorado at Denver, December 1994. <http://www-math.cudenver.edu/ccmreports/rep34.ps.gz>.
- [24] ———, *Convergence of algebraic multigrid based on smoothed aggregation*, UCD/CCM Report 126, Center for Computational Mathematics, University of Colorado at Denver, February 1998, <http://www-math.cudenver.edu/ccmreports/rep126.ps.gz>.